

CLAIMS

WHAT IS CLAIMED IS:

1. A method of recalling data objects stored on a plurality of media comprising:

receiving a plurality of requests to recall data objects, each data object being located on a particular one of said plurality of media;

creating a plurality of queues, wherein each one of said queues corresponds to one of said media;

placing said requests on the created queues, wherein each request is placed on the queue corresponding to the medium on which the requested data object is located;

activating a first of said queues, said first queue being associated with a first of said plurality of media; and

retrieving, from said first medium, the data objects requested on the first queue.

2. The method of claim 1, wherein each queued request indicates the offset from a predetermined location on a medium at which the requested data object is located, and wherein said placing act comprises:

organizing the requests on each queue in a first and a second sequence, each sequence comprising a set of requests whose offsets are monotonically increasing within the respective sequence.

3. The method of claim 1, further comprising:

activating a second of said queues, said second queue being associated with a second of said plurality of media, the second medium being different from said first medium; and

retrieving, from said second medium, the data objects requested on the second queue, the retrieval occurring at least in part concurrently with the retrieval of data objects from the first medium.

4. The method of claim 1, wherein each of said queue comprises a linked list of requests.

5. The method of claim 1, further comprising querying a database to determine the locations of the requested data objects.

6. The method of claim 5, wherein the location provided by said database comprises a media identifier and an offset into the identified medium.

7. The method of claim 1, wherein the act of creating a plurality of queues comprises timestamping each of the created queues, and wherein the act of activating a first of said plurality of queues comprises selecting a queue to be activated based on the timestamps of the respective queues.

8. The method of claim 1, said method being performed in a multi-threaded environment wherein plural threads execute concurrently, said method further comprising:

acquiring a lock prior to performing an action selected from the group consisting of:

placing a request on a queue; and
changing a queue from a non-active state to an active state;
said lock being acquirable by only one of said threads at a given time; and

releasing said lock subsequent to performing an action in said group.

9. The method of claim 1, wherein said method is performed in a computing environment comprising a hard disk, said method further comprising:

determining that said data objects are not located on said hard disk.

10. A computer-readable medium having computer-executable instructions to perform the method of claim 1.

11. In a computing environment having a first plurality of drives in which data objects to be retrieved are stored on a second plurality of media mountable on said drives, a method of obtaining data objects from said media comprising:

identifying a first data object located on a first of said media;

identifying a second data object located on a second of said media different from said first medium; and

concurrently using a first and a second of said drives to retrieve said first and second data objects from said first and second media.

12. The method of claim 11, further comprising:

creating a first queue corresponding to said first medium;

placing first data on said first queue, said first data being indicative of said first data object;

creating a second queue corresponding to said second medium;

and

placing second data on said second queue, said second data being indicative of said second data object.

13. The method of claim 12, wherein said first data comprises a location of said first data object on said first medium, and wherein said second data comprises a location of said second data object on said second medium.

14. The method of claim 11, further comprising:
 identifying a third data object located on a third of said media;
 waiting for said first or said second data object to be retrieved
 from their respective media; and
 retrieving said third data object from said third medium.

15. A computer-readable medium having computer-executable instructions to perform the method of claim 11.

~~16. A method of scheduling requests to recall data objects from a medium, said medium being mounted on a drive, said drive having a reading head, said method comprising:~~

receiving a new request to recall a data object from said medium,
said data object being located at a first offset along said medium;

identifying a queue comprising a plurality of requests to recall data from said medium, said requests having an order, each of said requests corresponding to a particular data object stored on said medium and indicating an offset into said medium at which the corresponding data object is located, said requests forming first and second sequences, said first sequence preceding said second sequence in said queue with respect to said order, wherein the offsets of the requests within said first sequence are a maximally monotonically increasing series;

determining the location of said reading head along said medium;

determining that the location of said reading head is beyond said first offset;

inserting said new request into said queue into a position that, with respect to said order, is subsequent to said first sequence.

17. The method of claim 16, wherein the offsets of the requests within said second sequence are a maximally monotonically increasing series.

18. The method of claim 17, wherein the position in said second sequence in which said new request is inserted is based on said first offset, such that the monotonically increasing nature of said second sequence is preserved following the insertion.

19. The method of claim 16, further comprising:

selecting, from among a plurality of queues, a particular queue into which to insert said new request, wherein each one of said plurality of queues corresponds to a different medium, the selection being based on the particular medium on which the data object is located.

20. A computer-readable medium having computer-executable instructions to perform the method of claim 16.

~~21. A system for retrieving data objects from a plurality of media comprising:~~

a queuing module which creates a plurality of queues corresponding to said plurality of media, said queuing module receiving requests to retrieve data objects from said plurality of media and queuing each of said requests on the queue corresponding to the medium on which the requested data object is located;

an activation module which selects queues for activation and activates the selected queues; and

a retrieval module which retrieves the items on one of the selected queues from the corresponding medium in the order in which the items are located on the queue.

22. The system of claim 21, further comprising:

~~a~~ a hard disk;

a file system which manages files on said hard disk, which stores information indicating which of said files have been migrated to said media, and which issues a request to said queuing module for requested files that have been migrated to said media.

23. The system of claim 21, wherein said queuing module timestamps each queue at the time that the queue is created, and wherein said activation module selects a queue for activation based on the timestamps of the created cues.

24. The system of claim 21, wherein each requested data object is located at an offset into the medium on which the data object is located, and wherein said queuing module includes logic which maintains the requests on each of said plurality of queues in first and second sequences, the offsets of the data objects requested in each of said first and second sequences comprising a maximally monotonically increasing series.

~~25.~~ A system for recalling data objects from a plurality of media comprising:

means for receiving requests to recall data;

means for creating a plurality of queues;

don't

means for placing the requests on said queues and
means for activating said queues; and
means for retrieving data requested on an activated

add B2

1. The first step is to identify the key components of the system. This involves understanding the hardware, software, and data involved.